

LA-UR-

*Approved for public release;
distribution is unlimited.*

Title:

Author(s):

Intended for:



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Data Intensive Science in the Department of Energy

October 2010

Abstract

In this white paper we provide a perspective on the opportunities and needs for data intensive science within the Department of Energy. In particular, we focus on two areas in which DOE's landscape is different from those of other organizations. First, DOE is a leader in the use of high performance computing for modeling and simulation, and these computations generate huge amounts of data to manage and analyze. Second, DOE maintains leading experimental facilities and these also produce prodigious quantities of data. In both of these areas, DOE's needs for data intensive science are distinct from those of other agencies, and we believe that these needs will likely not be adequately addressed without DOE investments.

1. Introduction

The growing ubiquity and volume of data is changing the landscape of commerce, science, and national security. Companies like Google use a form of data-intensive computing to organize information for the mass market. The Department of Defense maintains situational awareness through numerous sensors and data-intensive computing. NOAA and NASA generate vast quantities of earth observational data that, with the help of data-intensive computing, can be used to answer questions about planetary dynamics and ecosystems.

The diverse missions of the Department of Energy also rely heavily on the ability to extract meaning from data. Like scientists worldwide, DOE researchers rely on the Internet for knowledge discovery, and make use of bibliometrics and textual analysis of scientific publications. Like other large organizations, DOE defends its computer networks by analyzing large data streams in real time and identifying malicious content. In these and other areas, DOE's data-centric computing needs are similar to those of other organizations.

But there are two broad categories in which DOE's missions lead to unique data-centric computing needs, and these areas are the focus of this white paper.

First, DOE is internationally acknowledged as a leader in the use of advanced computing to simulate complex physical and engineering systems. DOE researchers routinely compute detailed models of time dependent, three-dimensional systems on the world's largest computers. These simulations generate enormous data sets that are difficult to extract and archive, let alone analyze. More comprehensive analysis of this data would help in the discovery and identification of unanticipated phenomena, and also help expose shortcomings in the simulation methodologies and software. As leadership-class computers continue to grow in size, the data analysis problems they inspire will be a major scientific challenge. These topics are explored more thoroughly in Section 2 below. They are also central themes in hybrid computational / experimental data analysis as discussed in Section 3.

Second, DOE manages some of the nation's most advanced experimental resources, and these facilities generate tremendous amounts of data. Datasets generated at DOE's advanced facilities today significantly outstrip current ability for analysis. As these facilities anticipate significant upgrades, the problem will increase tenfold and more. When datastreams are not optimally exploited, scientific discovery is delayed or missed. Also, real time analysis of the data as it is being generated would enable intelligent design and refinement of the experimental process. This would leverage the utilization of the facility and the quality of the science produced. This is not possible with current workflows and analytic capabilities. These challenges and opportunities are discussed in more detail in Section 4.

In Section 5 we identify common themes that emerge from the data-centric computing needs of scientific computing and experimentation. These lead to a set of recommendations in Section 6 that we believe will improve DOE's ability to advance science through data intensive computing.

There are a number of other relevant reports in the area of data-intensive science. A finding of the President's council of scientific advisors on science and technology (PCAST) report [1] of August 2007 states: "The data deluge represents an opportunity to advance U.S. leadership in science and technology, and harnessing it has become a national priority. More robust ... capabilities are needed to fully exploit large-scale data resources." The National Research Council of the National Academies produced a report that identified major challenges that hinder large-scale data integration [2]. The report addresses both issues of the massive scale of the data resources as well as issues of integrated access to multiple large-scale data collections. Critically important issues that address the quality and enduring legacy of scientific data are identified and discussed in a National Academy of Sciences report [3]. This report highlights three key data management issues: integrity, accessibility and stewardship. Our report also builds upon the work of the DOE data analysis community. In a 2007 DOE Advanced Scientific Computing Research (ASCR) office report [4], community members came together to identify how our data analysis and exploration tools need to change to process the massive datasets that will be generated by petascale and exascale supercomputing resources. The International Exascale Software Project [5] offers a roadmap of data management, analysis and visualization research and development efforts for successfully fielding exascale supercomputing resources.

This report is written by DOE scientists that deal with observational, experimental and simulation data who manage, analyze and understand data as part of their day-to-day mission. We believe our perspective supports and emphasizes the required focus on data issues identified in all the reports highlighted above.

2. Case Study: Scientific and Engineering Simulation Data

Computational modeling and simulation is central to numerous scientific and engineering domains. Computational results provide scientific insights that complement those of theory and experiment. The Department of Energy is an acknowledged leader in computational simulation, with state of the art facilities and simulation codes. Basic simulation data is often four dimensional (three spatial dimensions and time), but additional variable types, such as vector or tensor fields, multiple variables, multiple spatial scales, parameter studies, and uncertainty analysis can increase the dimensionality. There is no shortage of DOE applications that push the state of the art in high performance computing. Recent DOE exascale workshops [6] have reviewed the computational needs in fusion, nuclear energy, climate, national security, materials science, chemistry, biology, and more . In all of these scientific domains high-end simulations can generate copious amounts of data. Workflows and systems for managing and analyzing this data are already at the breaking point. And as computations grow in complexity and fidelity and run on larger computers the analysis of the data they generate will become more challenging still.

As just one example, nuclear energy researchers are contemplating the use of simulation to perform predictive science, using the results of simulations to understand the design space for new advanced nuclear reactors, to accelerate the design of such systems, and to optimize these designs. The codes developed to perform these calculations will leverage the most efficient algorithms possible and take advantage of the most powerful computers available. For these kinds of high-consequence simulations, it is essential that the codes be verified and validated, and that uncertainties be quantitatively understood. This requires an adaptive approach combining simulation and experimentation to resolve ambiguity and refine understanding.

A standard approach for analyzing and archiving computational data involves two distinct subsystems attached to the simulation platform via a large switch complex. A separate subsystem consisting of many storage servers backed by enterprise storage provides the raw storage for scientific data. An analysis cluster, also attached to the switch complex, is available to analyze archived data. But this approach requires that data be moved off the simulation machine prior to analysis. As the gap between compute performance and I/O performance grows, data generation rates are increasingly problematic, and scientists are struggling to reduce their data output in order to minimize the cost to write data at runtime, and the cost to analyze the data subsequently.

DOE's detailed simulations require the fastest and largest supercomputers available today and in the future. The largest supercomputer today is in the petascale class; current planning is focused on exascale class machines. The chip, memory, and networking technology trends that drive exascale architectures suggest that limiting factors include power and therefore data movement (a power intensive operation). Since data movement dominates the computational costs at exascale, rethinking the entire data analysis process is required.

An in-depth data triage needs to occur while the data resides in the memory of the supercomputing platform. This will require fundamental new massively parallel data reduction algorithms in order to effectively analyze these massive datasets. Example approaches include statistical and compression techniques that downsample the data before transfer. These approaches can provide multiple levels of resolution that highlight areas of interest to the scientist. A key approach is the integration of science-driven feature extraction algorithms that identify higher-order structures from the simulation results. These structures are information-rich but smaller in size than the massive datasets they are derived from.

An important aspect of this process is preserving the scientific integrity of the analyzed data. Recording the provenance of data, (i.e. how it was created and processed) as well as understanding if and when bias is introduced during the analysis process supports reproducible science. Data management processes that support effective data archiving, curation, and sharing need to be part of the analysis workflow from beginning to end.

An additional challenge comes from the increasing analytical complexity required for further scientific insight. Good methods exist for visualizing several variables in a subset of the simulation output. However, as simulations gain complexity and fidelity, it is necessary to explore the output more holistically to search for unexpected behavior. This is essential to the validation process, and also to the scientific discovery process. Exploratory analysis requires efficient techniques for accessing disparate portions of the simulation output, and a high degree of user interactivity. However, methods for exploring fundamental issues like uncertainty quantification remain in their infancy. These kinds of advanced analysis capabilities are quite difficult to provide for very large data sets.

I/O speeds are only one of several technological impediments to the analysis of simulation data. As illustrated in Figure 1, the historical rate of improvement in disk performance is dwarfed by the historical rate of growth in the computational power of HPC systems. This results in an unsustainable mismatch between the rate at which computational data can be generated, and the rate at which it can be archived. As a result, most computational data is never stored, and only periodic snapshots are available for analysis.

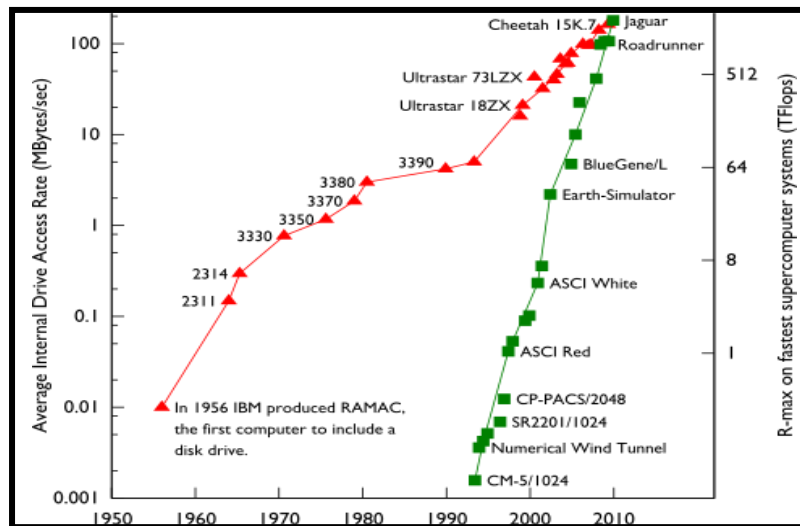


Figure 1: The rate of performance of improvement in disks is much lower than the rate of improvement in compute systems, driving the need for ever-larger disk counts in each successive HPC system deployment

simply to keep pace. From the cost, reliability, and power perspectives, this approach is not sustainable. (Thanks to R. Freitas of IBM Almaden Research for providing some of this data.)

One encouraging technology development is the rapid maturation of solid state storage. It is not yet clear how best to make use of these devices – whether they should be considered as an alternative to disks, or whether the entire system architecture should be reconsidered. But their ability to support large random reads will make them an important component of the data intensive toolkit.

As simulations continue to grow in size and fidelity, they will increasingly be used for critical decision-making. This will require the computing community to exercise diligence in the management of software and of data.

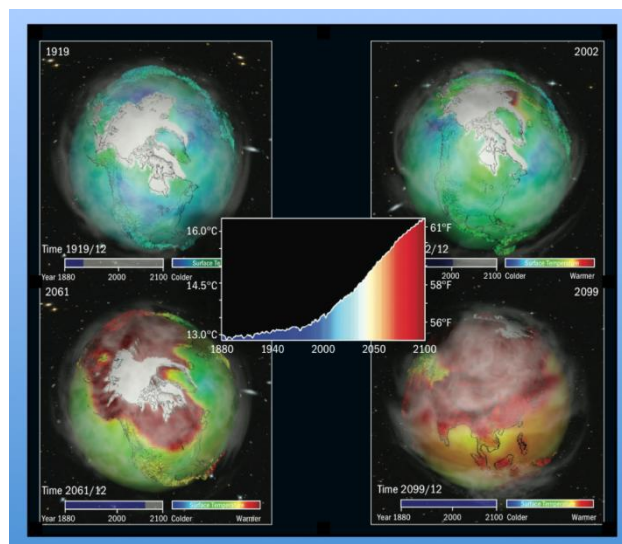
Summary:

- Current approaches for analyzing and visualizing scientific simulations that involve the movement of data will become difficult or impossible as we transition to exascale. Therefore significant reductions in the amount of data written to storage systems will be required with larger datasets.
- The exascale roadmaps indicate that raw computing power will grow much more rapidly than bandwidth, I/O, or storage capacity. These changes will further exacerbate the challenges associated with the analysis of computational data. Fundamentally new paradigms for data-oriented system design and workflow are urgently needed.
- In addition to system changes, new scientific analysis techniques for verification and validation, uncertainty quantification, multi-scale physics, data integrity and statistical results will need to be developed. As the sophistication of simulations increases so must the sophistication of analysis methods that help derive insight from them.

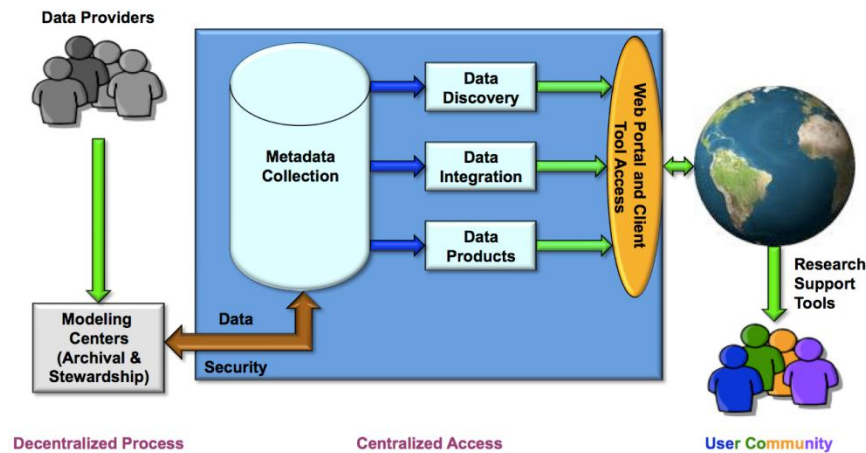
3. Case Study: Climate Data

Climate scientists and climate research data play a crucial role in understanding our planet and in shaping the political responses to climate change. Observational data show that over the last century, Earth's global mean temperature has risen more than 1°C. Over the next century, simulated data predict another 3 to 6°C increase due to the burning of fossil fuels and other human activities [7].

To help better understand the consequences of global warming, scientists are looking at a diverse set of data sources obtained from model simulation experiments, remote sensors (e.g., radar and satellites), and mobile as well as in situ (e.g., a radiosonde measuring a parcel of air or an anemometer measuring wind) observational platforms. Producing hundreds of petabytes (PB, where 1 PB is 1×10^{15} bytes) of data, these multiple data sources must be tightly integrated into a “smart data infrastructure system” to help scientists project the impacts of future climate change and define options for mitigating and adapting to that change.



Climate science relies upon facilities that offer access to computational resources, data storage and movement, workflows and provenance, and analysis tools—to name a few. Simply stated, Climate change research is not only a scientific challenge of the first order, but also a major technological and infrastructure challenge.



In response to these challenges, climate scientists and computational scientists are collaborating worldwide to assemble the largest-ever collection of simulation and observation data sets for the Fifth Phase of the Coupled Model Intercomparison Project (CMIP5). CMIP5 is expected to provide results that will be fundamental to the 2013 scientific Fifth Assessment Report (AR5) of climate science by the Intergovernmental Panel on Climate Change (IPCC). By comparison, the CMIP5 data archive will dwarf that of its predecessor, the Third Phase of the Coupled Model Intercomparison Project (CMIP3). Released for climate study in 2004, the CMIP3 archive provided 35 terabytes of data and was used for the 2007 Nobel Prize-winning IPCC Fourth Assessment Report (AR4). For CMIP5, modeling groups will generate tens of petabytes of data, and the ~2-PB subset of data expected to be of highest interest to researchers will be replicated at several data centers around the world. As with CMIP3, the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at Lawrence Livermore National Laboratory will supervise the distributed CMIP5 data archive and oversee the effort to provide access to this valuable collection of output from model simulations and observations.

In this challenging undertaking, PCMDI has organized partnerships with global data centers funded to assist with CMIP5 data retrieval and dissemination to create an internationally distributed data archival and retrieval federation. Their mission is to provide climate researchers worldwide with a science gateway to access the data, information, models, analysis tools, and computational capabilities required to evaluate extreme-scale data sets. Its stated goals are to (1) make data more useful to climate researchers by developing collaborative technology that enhances data usability; (2) meet the specific needs of national and international climate projects for distributed databases, data access, and data movement; (3) provide a universal and secure Web-based data access portal for broad-based multi-model data collections; and (4) provide a wide range of climate data-analysis tools and diagnostic methods to international climate centers and U.S. government agencies. To this end, the partners are working to integrate all important climate data sets—from climate simulations to observations—using distributed storage management, remote high-performance units, high-bandwidth wide-area networks, and user desktop platforms in a collaborative problem-solving environment.

In building the distributed climate archives, modeling groups provide data through an ESG Data Node. Actual data holdings reside on a large number of federated ESG Data Nodes. The nodes host those data and the metadata services needed to publish data onto ESG and execute data-product requests through an ESG Gateway. The ESG Gateways, which act as brokers handling data, process requests to serve specific user communities. Services deployed on a gateway include the user interface for searching and browsing metadata, for requesting data products (including analysis and visualization tools), and for orchestrating complex workflows. These large-scale computations will frequently involve numerous resources spread throughout the modeling and observational climate communities. Aggregations of data and computational methods are varied; thus, needed resources must be positioned strategically throughout the global community to facilitate research to understand the complex nature of important climate processes.

Based on the growth rates established in the DOE report, Scientific Grand Challenges: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale [8], data from the climate community will reach hundreds of exabytes by 2020. To be ready for this expected onslaught, the ESG federated enterprise system is in place and awaiting the first set of CMIP5 simulation and observational data. Besides the CMIP5 archive, NASA has also partnered with the ESG-CET team to disseminate NASA's satellite data (e.g., Atmospheric Infrared Sounder, Microwave Limb Sounder, Cloudsat) and observational data archives. In addition, the National Oceanic and Atmospheric Administration has partnered with the ESG-CET team to make the holdings of the National Climate Data Center, the world's largest active archive of weather data, available to climate researchers. With more than 20,000 registered users and over 1 PB of data downloaded by the community, ESG-CET is quickly expanding to encompass greater and greater international climate data domains, paving the way for a truly global network of climate data and services whose goal is to make climate data easily accessible and easy to understand.

Climate science data sets are growing at an extremely fast rate [9]. Table 1 shows key science drivers and anticipated storage and network requirements for the future climate community.

Table 1: Climate community anticipated requirements

Feature	Key Science Drivers		Anticipated Computing and Storage Requirements	Anticipated Network Requirements	
Time Frame	Science Instruments and Facilities	Process of Science	Computing and Storage Resources	Local-Area Network Bandwidth and Services	Wide-Area Network Bandwidth and Services
Near-term (0–2 years)	<ul style="list-style-type: none"> • ESG federated Data Node • ESG Gateways • Linux front-end with 10×10-gigabit-per-second (Gbps) connections 	<ul style="list-style-type: none"> • CMIP5/AR5 data access • NASA satellite observations • NOAA Numerical Weather Prediction 	<ul style="list-style-type: none"> • Large supercomputer center • Petaflops computing • Tens of PB rotating file system • Broad user base • Tens of PB high-performance storage system (HPSS) • Remote analysis 	<ul style="list-style-type: none"> • 1- to 10-gigabit Ethernet (GigE) 	<ul style="list-style-type: none"> • 10×10-Gbps network connection
Mid-term (2–5 years)	<ul style="list-style-type: none"> • Extend ESG Data Nodes to other climate archives • 100-Gbps connection 	<ul style="list-style-type: none"> • IPCC AR5 Report • Community Earth System Model (CESM), Geophysical Fluid Dynamics Laboratory, etc. model development • Climate Science for a Sustainable Energy Future testbed and uncertainty quantification activities 	<ul style="list-style-type: none"> • Cloud computing facilities • Large supercomputing center • Exaflops computing • Hundreds of PB rotating file system • Hundreds of PB HPSS • Geographic information system integration 	<ul style="list-style-type: none"> • Expand use of 10 GigE 	<ul style="list-style-type: none"> • 100+ Gbps network connection
Long-term (5+ years)	<ul style="list-style-type: none"> • Extend ESG Data Nodes to other science domains • Network connections to 10×100 Gbps 	<ul style="list-style-type: none"> • CMIP6 (IPCC AR6) • Other climate simulation and observational data • Climate and economic model 	<ul style="list-style-type: none"> • Cloud computing facilities • Large supercomputing center • Multiple exaflops computing • Exabyte rotating file system • Exabyte HPSS 	<ul style="list-style-type: none"> • Expand use of 100 GigE 	<ul style="list-style-type: none"> • 10×100-Gbps network connection

Summary:

- In order to understand the impacts of climate, climate simulations and observations studies will be

conducted. These simulations and observations will produce hundreds of petabytes of data that needs to be integrated, curated and available for analysis.

- In order to effectively handle these massive datasets, an advanced networking and analysis infrastructure must be researched, developed and deployed.

4. Case Study: X-Ray and Neutron Scattering Data

Over the past two decades, scientific research at intense X-ray and neutron sources in the United States has had a major impact on the way we understand the world around us. Physics, materials science, chemistry, biology, geosciences, medicine, engineering and the energy sciences have advanced significantly based on discoveries at these national facilities. While measurement capabilities are uncovering previously inaccessible information, and the future is expected to bring these at an ever increasing rate, data challenges are limiting our ability to fully interpret the experiments. In particular, while experimental capabilities are improving rapidly, robust capabilities for handling the data and for the corresponding theory, modeling and simulation for leveraging and fully benefiting from these results lag far behind.

As a concrete example, consider the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory that provides the most intense pulsed neutron beams in the world for scientific research and industrial development. Neutron scattering research has been used to develop many of the advanced materials that have impact in medicine, electronics, aerospace, and other areas of everyday life. As the SNS develops to its full realization, its software analysis and data handling platforms will require enormous effort to reach commensurability with the experimental capabilities.

Another example is the Advanced Photon Source (APS), which provides the brightest storage-ring generated X-ray beams in the Western Hemisphere. With the new APS upgrade project, capabilities will reach nanometers of spatial resolution for imaging and picoseconds of timing resolution for dynamics research. The two major APS themes of “mastering hierarchical structures through imaging”, and “real materials under real conditions in real time” will ensure US leadership in X-ray based research and development but will require implementation of significant computational resources. Here as well, software analysis and data handling platforms will require a ten-fold expansion to reach commensurability with the upgraded experimental capabilities.

Detailed SNS data streams entail space (x,y) and time coordinates for every detected neutron. The data can be streamed for processing during an experiment, or concatenated to produce an “event list” data file for post measurement analysis. Often experiments are performed varying one or more parametric values such as temperature, pressure, magnetic field, gas/liquid flow rates, or by rotation angle – thus raw data starts at 3D and ranges to as many dimensions as parametric variables are studied. These 3D data sets typically range from 10 to 100 GBytes in size, and to process a raw data set may require the simultaneous utilization of up to 6 calibration and correction data sets while also carrying along the proper experimental error estimates. Thus the composite raw data sets can approach a Terabyte of data to process per measurement, per instrument – before adding parametric variations.

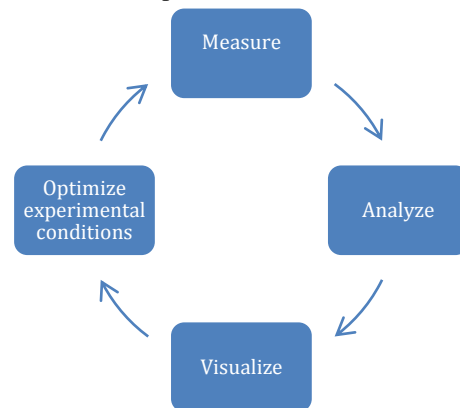
When using multiple measurements for tomographic reconstruction of a sample, the data and computational requirements quickly become daunting. Measurements from 100 angles will result in at least 100 Terabytes of data that need to be managed and processed. At the APS, typical measurements today make use of 720 angles, and real time implementations are planned for the future. The tomographic reconstruction algorithms use an optimization technique that typically requires $O(1000)$ iterations to converge. The total computational requirements are on the order of 10^{18} operations. And these operations involve complex data access patterns, which results in poor utilization of traditional high performance computers. To avoid these difficulties, the data is typically reduced prior to analysis, inevitably losing information. Improvements in data processing capabilities are essential to make maximal use of state-of-the-art experimental facilities.

Today it is well understood that single measurements alone rarely provide unambiguous results to complex questions. New science and basic understanding are made possible by software that coordinates analyses across experimental techniques and user facilities. For example, the frontier in materials characterization exists where a variety of physical measurements, informed by theory, modeling and simulations, are used to

solve the important and difficult problems. Users require experiments that can be done in an environment capable of accumulating, managing, and analyzing the diverse data they will collect. While the rewards are anticipated to be great, such systems today are beyond the horizon. Advancement towards this frontier will require a new generation of software for data analysis and modeling.

Starting with the simplest issues, software for data analysis, modeling and simulation at X-ray and neutron facilities begins as soon as data reduction and instrumental corrections are complete. In this area, a primary need is the ability to do on-demand real time analysis of data while it is being collected. This basic requirement ensures that the intended data is being measured and experimenters are not following blind alleys. Odd as it may sound, most of the data acquired at the large facilities today are not subject to quality evaluation and interpretation until after the measurements are complete.

The large data streams associated with current and future experimental facilities are driving a change in the way experimental data is managed. It once was possible for users to process their data on their own computing resources, but this is no longer a viable solution. Instead, data management and processing must be an integral aspect of the experimental facility. SNS researchers have access to computing facilities at SNS and at ORNL, but these are primarily used for data reduction and the resulting reduced data is left for users to analyze at their home facilities.



However a new scientific vision is emerging which more tightly couples experimentation, data analysis and computational modeling. In this vision, computation will be used to provide scientific inference to users as they perform their experiments. Experiments will engage a feedback loop to measure, analyze, visualize, and then optimize for the next measurement iteration. This vision is being realized on a modest scale today, but its full realization will require faster analysis and simulation capabilities to support real-time decision-making.

Indeed, one step before this, if predictive modeling and simulation were available in advance of the measurements, it would expand greatly our insights into the experiments. In most areas, this capability is not available at the facilities. In these very fundamental areas computational scattering science is a most promising opportunity for expanding our insights into the leading edge science that is being performed [10]. These opportunities can only be realized by dramatic changes in the way data is managed and analyzed at experimental facilities. The workflows and systems designed for traditional computational science will need to be redesigned for hybrid experimental/computational science.

Summary:

- Experimental facilities are facing massive data challenges. Issues of data integrity and quality are important throughout the required data reduction process from sensor output to storage.
- Tightly coupled experimentation, modeling/simulation and analysis are the way cutting edge science is conducted today. Data analysis, comparison and archiving are at the heart of this scientific process. To facilitate successful science, data issues including massive sizes, real-time analysis, data integrity and advanced analysis issues must be addressed.

5. Overall DOE data intensive landscape

The case studies sketched above differ in many details, but they share underlying features that are common to other scientific data problems. In this section, we explore some of these common needs and extract some broader lessons about DOE's challenges.

Data Issues

First and foremost, data is important. Data is the source of all scientific insight that emerges from a computation or an experiment. Given the investment required to obtain that data, it is essential that data be managed in a disciplined and mindful manner to preserve its integrity and its value for anticipated and unanticipated uses. Careful data management requires well-designed policies, architectures, and workflows. Appropriate hardware and software tools are required to collect, move, and store data. Integrity can be enhanced through access controls, cryptographic protocols, and redundancy. Generally speaking, experimental communities have addressed these challenges more holistically than computational scientists. Commercial off-the-shelf products are available to meet many of these needs, particularly if the data volumes are modest.

The analysis of scientific data is often highly specialized and depends upon the characteristics of the data and the precise questions being addressed. Simple scientific analytics may consist of summary statistics or selections of subsets of data. These are easily addressed by traditional data methods like database queries or MapReduce, or via custom tools written by domain scientists. However, more complex analytical problems can be considerably more challenging. For example, data from a fluid simulation might be analyzed to characterize recirculation regions. A new genomic fragment might be compared to a library of existing sequences. A series of computations might be studied to quantify margins of uncertainty. These kinds of specialized analytics are not well addressed by general-purpose tools. The functionality of databases and MapReduce approaches is sometimes too limited to easily express the analytical need, and these approaches can impose large inefficiencies. Instead, the scientific community has traditionally written customized algorithms and software for specific data analysis challenges. The complexity of these specialized analyses is compounded by the challenge of working with large scale data. As a result, analysis tends to be focused on answering very specific questions, and not on open-ended exploration. Potential scientific discoveries are missed.

Although a computer can answer specific questions about a data set, much richer understanding can emerge by allowing a scientist to explore the data interactively. Considerable work has gone into building interactive visualization environments for simulation data. This data has very special structure – often an underlying mesh or grid in three dimensions over time – that allows for efficient and intuitive exploration.

Data that does not have an underlying physical geometry can be much more difficult to explore. Examples include biological pathway analysis, uncertainty characterization from ensembles of simulations, and cyber defense. For these kinds of applications, novel data abstractions and representations are needed to communicate to a user. The choice of an abstraction and its visualization imposes a mental model on the user that can be highly informative, but also deceptive. Spurious correlations and outliers can be an artifact of the human-computer interface, and not germane to the data. The development of environments for richly exploring complex data is a very active area of research.

As data sources become larger and more complex, the challenges become even more acute. Even today, in both experimental and computational settings the vast majority of data is never analyzed in any depth by humans or by computers. Instead, a scientist typically studies data summaries or small subsets of the overall data to confirm or refute a hypothesis. All of the rest of the data is generated, transferred, and archived, but ignored. It is difficult to know how much deeper understanding is lurking in the unexamined and unanalyzed data. In the future, computers will need to play a growing role in automatically identifying interesting phenomena and bringing them to the attention of scientists. Instead of just answering well-defined questions about the data, advanced analysis systems should be able to find interesting patterns, trends, and anomalies. To do this in a meaningful way will require advances in statistics, machine learning, information theory, and other areas of mathematics and computer science.

Software Issues

There is need for much faster analysis of data streams. If scientific data could be studied in situ, it would enable interactive steering for both computational and experimental science. In both domains, expensive resources are given a prescribed task and set loose with minimal oversight. In both worlds many tasks fail to deliver as planned and no useful data is produced. An alternative would be to actively steer the computation or the experiment based upon the characteristics of the data being generated. Computations that have become non-physical could be ended early, or regions of parameter space that are known to cause instabilities could be steered around. For experimental studies, parameters could be adapted on the fly to ensure the quality and the relevance of the data being collected, and experiments could be redirected as soon as the necessary data has been gathered. In both domains the result would be higher quality science and improved utilization of precious resources.

Interactive steering requires real time analysis of data to either self-monitor or to support a human decision maker. Data analysis and transmission would need to be designed into the workflow to a degree that it is generally not today. Sufficient and appropriate computational resources for analyzing data in-situ would need to be part of the system architecture.

Another important scientific opportunity is insight that could arise from closer coupling between different experiments or between experimentation and simulation. For example, using first order difference methods to combine neutron and x-ray structure factors, atomic and nanoscale ordering can be observed in unprecedented detail compared to either method alone. Each experimental and computational method provides distinct insight, but their combination offers information that is otherwise inaccessible. Realizing this vision will require better methods for managing and aligning multiple data sets, and techniques for analyzing them simultaneously.

Architectural Issues

All of these data-intensive computing opportunities require advanced computational capabilities. Unfortunately, we believe the current high performance computing ecosystem is not a good match for the needs of many DOE data intensive applications. Data intensive applications can require advanced computing to perform analytics for large and complex data sets, or they may require answers to be computed quickly for real time applications like cyber security or experimental steering. In either case, the needs of data intensive applications are distinct from those of the modeling and simulation applications that have driven traditional high performance computing. We envision the need for advances across a wide spectrum of the computing landscape.

Data intensive computing will place high demands on hardware. Data intensive applications require less computing power than modeling and simulation applications, but they place more demand on all layers of the memory and communication systems. Improvements in the effective performance and power efficiency of memory are critical. Networks are needed that have higher bandwidth and lower latency, particularly for small messages. Current processor architectures rely on high spatial and temporal locality to make effective use of memory hierarchies, but many data intensive applications have poor locality. Architectures that exploit massively multithreading may be able to achieve much higher performance for data applications while consuming much less power. It is worth exploring even more radical memory-centric designs in which memory requests computing services instead of processors requesting data from dumb memory.

Systems issues for data intensive applications are also quite different from those associated with modeling and simulation. The balanced machine for data applications will require more memory and bandwidth, and comparatively less compute power. Data applications can benefit greatly from a shared address space and programming models that support fine-grained parallelism instead of just bulk synchronous parallelism. Many data intensive applications benefit from a close interaction with a human decision maker. Having a human in the loop requires fresh thought about how a high performance computer is managed and scheduled.

Although the needs of data intensive computing are distinct from those of traditional HPC applications, they share some the key technological underpinnings. More effective and power-efficient memory systems are critical to future generations of high performance computers. Memory power and performance will be

the limiting factors in the drive towards exascale computing. The advances in memory systems enabled by investments in exascale computing will be hugely important for data intensive applications. The growing complexity of scientific applications coupled with the vast computational power of many-core nodes means that future modeling and simulation codes will not be limited by computational speed. Instead, memory and network performance will be the key constraint on performance. This will require improvements to system and network technologies that will also benefit data intensive applications. In addition, the use of many-core nodes will require new programming models and paradigms that can take advantage of shared memory. These new approaches may simplify the development of data intensive applications.

Summary:

- DOE missions require the analysis of diverse types and sizes of data including scientific and sensor-based data. Issues of data integrity, provenance and ease of analysis all need to be dealt with carefully as this data is the source of our primary product: scientific understanding and insight.
- Integrated data analysis approaches that reduce data movement and bridge the gap between experiments, simulations and analysis are critical to reducing the problem of analyzing massive data. This tightly coupled approach will improve the scientific process by reducing barriers between these different sources of scientific knowledge.
- Traditional HPC architectures are not designed to efficiently process data-intensive workloads. A multi-pronged approach that includes the integration of data-intensive features into HPC architectures as well as designing custom data-intensive architectures from scratch will help us efficiently handle massive datasets.

6. Recommendations

As a mission-driven agency DOE is on the front lines of addressing data-intensive science needs. As such, effective solutions must be an integral part of the science endeavor that is underway. One reason for this is obvious; generically developed solutions are difficult to interface to and incorporate into a project workflow. More profoundly, however, is the bandwidth limitation imposed by today's and tomorrow's gigascale networks. Data from petascale and exascale simulations and experiments will overwhelm our networks. Analysis must occur immediately, at the data source and in a scalable manner. This is why promising technology schemes such as remote-site cloud computing are unlikely to address the looming challenge of massive datasets.

Current limitations in data intensive computing are significant impediments to the advancement of science in the Department of Energy. On current trends, these problems will get worse in coming years as experimental and computational data streams increase in size and complexity. We have three broad recommendations to address these challenges:

(1) Data management and analysis should become a central component of large experimental or computational programs. Data issues should be considered as an integral component of the system design, the hardware investments, the workflow design, and the staffing of facilities. Since data is so central to our mission we believe there should be a recognized and supported data-intensive focused program that goes beyond specific DOE programs. A key part of this data-intensive program is the deployment of research and development to support mission-driven programs in both simulation and experimental programs. The impact of data-intensive technology on both simulation and experimental programs should be used as a measure of the success of any data-intensive computing program.

(2) Fundamental research investments in data-intensive computing must be made to maximize the scientific insights obtained from DOE's experimental and computational programs. These investments should span the full computing ecosystem, leveraging existing DOE expertise and focus where possible. Areas in need of advances include the following.

- (i) Hardware including memory systems, interconnects, and processor design.
- (ii) System architecture including programming models and workflows for data intensive problems.
- (iii) Integrated approaches that reduce data movement by compressing and analyzing data as it generated.
- (iv) Algorithms including scalable statistical and machine learning techniques for scientific data.
- (v) Human-computer interfaces including visualization paradigms and human-in-the-loop supercomputing.

(3) A successful R&D program is dependent upon the quality of the available research staff. Data-intensive computing is a difficult, long-term problem for DOE. As part of an integrated plan, we believe an education program, starting at the undergraduate level that includes classes, internships and university collaborations is required. Creating the next-generation workforce of data/knowledge will support DOE's data needs well into the future.

References

- [1] J. Marburger and E. Kvamme, Leadership Under Challenge: Information Technology R&D in a Competitive World -- an Assessment of the Federal Networking and Information Technology R&D Program, President's Council of Advisors on Science and Technology, August 2007.
- [2] S. Weidman and T. Arrison, Steps Toward Large-Scale Data Integration in the Sciences, National Research Council of the National Academies, 2010.
- [3] T. Arrison and D. Stine, Committee on Ensuring the Utility and Integrity of Research Data in a Digital Age, Ensuring the Integrity, Accessibility, and Stewardship of Research Data in the Digital Age, National Research Council of the National Academies, 2009.
- [4] C. Johnson and R. Ross, Visualization and Knowledge Discovery: Report from the DOE/ASCR Workshop on Visual Analysis and Data Exploration at Extreme Scale, June 2007.
- [5] International Exascale Software Project, Roadmap 1.0, May 2010.
- [6] DOE Exascale Reports, <http://www.science.doe.gov/ascr/ProgramDocuments/ProgDocs.html>.
- [7] IPCC (Intergovernmental Panel on Climate Change), 2007: Climate Change 2007. www.ipcc.ch.
- [8] Washington et al. 2008, Scientific Grand Challenges: Challenges in Climate Change Science and the Role of Computing at the Extreme Scale, 2008. <http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ClimateReport.pdf>
- [9] ASCR Science Requirements, 2009: Office of Advanced Scientific Computing Research Network Requirements Workshop 2009. <http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ASCR-Net-Req-Workshop-2009-Final-Report.pdf>.
- [10] Report from the Workshop on Computational Scattering Science, 6 – 9 July 2010, Argonne National Laboratory. Organized by Brent Fultz, Kenneth Herwig and Gabrielle Long.

Authors (in alphabetical order)

Jim Ahrens, Los Alamos National Laboratory¹

Bruce Hendrickson, Sandia National Laboratories²

Gabrielle Long, Argonne National Laboratory³

Steve Miller, Oak Ridge National Laboratory⁴

Robert Ross, Argonne National Laboratory³

Dean Williams, Lawrence Livermore National Laboratory⁵

¹Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.

²Sandia is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

³Long and Ross were supported by the U.S. Department of Energy Office of Science under Contract No. DE-AC02-06CH11357.

⁴Oak Ridge National Laboratory is managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR22724.

⁵This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.